

The Claims:

1. (Currently Amended) A method comprising:
determining that one of a plurality of nodes has failed, each node comprising a switching fabric integrated ~~onto a board~~ to a card and ~~one or more~~ at least two processors integrated ~~onto the board~~ to the card;
removing the failed node from a virtual list of nodes, the virtual list comprising one logical entry for each of the plurality of nodes;
determining that at least a portion of a job was being executed on the failed node;
terminating at least the portion of the job;
determining that the job was associated with a subset of the plurality of nodes; and
deallocating the subset of nodes from the job.
- 2-3 (Canceled)
4. (Previously Presented) The method of Claim 1, each entry of the virtual list comprising a node status and the method further comprising changing the status of each of the subset of nodes to "available."
5. (Previously Presented) The method of Claim 1, further comprising:
determining dimensions of the terminated job based on one or more job parameters and an associated policy;
dynamically allocating a second subset of the plurality of nodes to the terminated job based on the determined dimensions; and
executing the terminated job on the allocated second subset.
6. (Original) The method of Claim 5, the second subset comprising a substantially similar set of nodes to the first subset.

7. (Previously Presented) The method of Claim 5, wherein dynamically allocating the second subset comprises:

determining an optimum subset of nodes from a topology of unallocated nodes; and
allocating the optimum subset.

8. (Previously Presented) The method of Claim 1, further comprising:
locating a replacement node for the failed node; and
updating the logical entry of the failed node with information on the replacement node.

9. (Previously Presented) The method of Claim 1, wherein determining that one of the plurality of nodes has failed comprises determining that a repeating communication has not been received from the failed node.

10. (Previously Presented) The method of Claim 1, wherein determining that one of the plurality of nodes has failed comprises determining through polling that one of the plurality of nodes has failed.

11. (Currently Amended) Software encoded in one or more computer-readable tangible media and when executed operable to:

determine that one of a plurality of nodes has failed, each node comprising a switching fabric integrated ~~onto a board~~ to a card and ~~one or more~~ at least two processors integrated ~~onto the board~~ to the card;

remove the failed node from a virtual list of nodes, the virtual list comprising one logical entry for each of the plurality of nodes;

determine that at least a portion of an job was being executed on the failed node;

terminate at least the portion of the job;

determine that the job was associated with a subset of the plurality of nodes; and

deallocate the subset of nodes from the job.

12-13 (Canceled)

14. (Previously Presented) The software of Claim 11, each entry of the virtual list comprising a node status and the software further operable to change the status of each of the subset of nodes to "available."

15. (Previously Presented) The software of Claim 11, further operable to:
determine dimensions of the terminated job based on one or more job parameters and an associated policy;
dynamically allocate a second subset of the plurality of nodes to the terminated job based on the determined dimensions; and
execute the terminated job on the allocated second subset.

16. (Original) The software of Claim 15, the second subset comprising a substantially similar set of nodes to the first subset.

17. (Previously Presented) The software of Claim 15, wherein the software is operable to dynamically allocate the second subset comprises software operable to:
determine an optimum subset of nodes from a topology of unallocated nodes; and
allocate the optimum subset.

18. (Previously Presented) The software of Claim 11, further operable to:
locate a replacement node for the failed node; and
update the logical entry of the failed node with information on the replacement node.

19. (Previously Presented) The software of Claim 11, wherein the software being operable to determine that one of the plurality of nodes has failed comprises the software being operable to determine that a repeating communication has not been received from the failed node.

20. (Previously Presented) The software of Claim 11, wherein the software being operable to determine that one of the plurality of nodes has failed comprises the software being operable to determine through polling that one of the plurality of nodes has failed.

21. (Currently Amended) A system comprising:
a plurality of nodes, each node comprising a switching fabric integrated ~~onto a board~~ to a card and ~~one or more~~ at least two processors integrated ~~onto the board~~ to the card; and
a management node operable to:
determine that one of the plurality of nodes has failed;
remove the failed node from a virtual list of nodes, the virtual list comprising one logical entry for each of the plurality of nodes;
determine that at least a portion of an job was being executed on the failed node;
terminate at least the portion of the job;
determine that the job was associated with a subset of the plurality of nodes; and
deallocate the subset of nodes from the job.

22-23 (Canceled)

24. (Previously Presented) The system of Claim 21, each entry of the virtual list comprising a node status and the management node further operable to change the status of each of the subset of nodes to "available."

25. (Previously Presented) The system of Claim 21, the management node being further operable to:
determine dimensions of the terminated job based on one or more job parameters and an associated policy;
dynamically allocate a second subset of the plurality of nodes to the terminated job based on the determined dimensions; and
execute the terminated job on the allocated second subset.

26. (Original) The system of Claim 25, the second subset comprising a substantially similar set of nodes to the first subset.

27. (Previously Presented) The system of Claim 25, wherein the management node being operable to dynamically allocate the second subset comprises the management node being operable to:

determine an optimum subset of nodes from a topology of unallocated nodes; and
allocate the optimum subset.

28. (Previously Presented) The system of Claim 21, the management node being further operable to:

locate a replacement node for the failed node; and
update the logical entry of the failed node with information on the replacement node.

29. (Previously Presented) The system of Claim 21, wherein the management node being operable to determine that one of the plurality of nodes has failed comprises the management node being operable to determine that a repeating communication has not been received from the failed node.

30. (Previously Presented) The system of Claim 21, wherein the management node is operable to determine through polling that one of the plurality of nodes has failed.

31. (Currently Amended) A method comprising:

determining that one of a plurality of ~~nodes~~ nodes has failed, each node comprising:

at least two first processors operable to communicate with each other via a direct link between them, the first processors integrated ~~onto a single first motherboard~~ to a first card; and

a first switch integrated ~~onto the single first motherboard~~ to the first card, the first processors communicably coupled to the first switch, the first switch operable to communicably couple the first processors to at least six second ~~motherboards~~ cards each comprising at least two second processors integrated ~~onto the second motherboard~~ to the second card and a second switch integrated ~~onto the second motherboard~~ to the second card operable to communicably couple the second processors to the first ~~motherboard~~ card and at least five third ~~motherboards~~ cards each comprising at least two third processors integrated ~~onto the third motherboard~~ to the third card and a third switch integrated ~~onto the third motherboard~~ to the third card;

the first processors being operable to communicate with particular second processors on a particular second ~~motherboard~~ card via the first switch and the second switch on the particular second ~~motherboard~~ card;

the first processors being operable to communicate with particular third processors on a particular third ~~motherboard~~ card via the first switch, a particular second switch on a particular second ~~motherboard~~ card between the first ~~motherboard~~ card and the particular third ~~motherboard~~ card, and the third switch on the particular third ~~motherboard~~ card without communicating via either second processor on the particular second card;

removing the failed node from a virtual list of nodes, the virtual list comprising one logical entry for each of the plurality of nodes;

determining that at least a portion of a job was being executed on the failed node;

terminating at least the portion of the job;

determining that the job was associated with a subset of the plurality of nodes; and

deallocating the subset of nodes from the job.